

06/14/99
jc520 U.S. PTO

Case/Docket No. TRANS11
Express Mail No.EJ738560499U

A
jc526 U.S. PTO
09/332338
06/14/99

THE HONORABLE COMMISSIONER OF PATENTS AND TRADEMARKS
Washington, D.C. 20231

Sir:

Transmitted herewith for filing is the patent application of:

Inventors: Guillermo J. Rozas, Godfrey P. D'Souza, Charles R. Price, and Paul S. Serris

For: METHOD AND APPARATUS FOR ENHANCING SCHEDULING IN AN
ADVANCED MICROPROCESSOR

Enclosed are:

- XXX Five (5) sheet(s) of Formal Drawing(s) including seven (7) figures.
XXX An Assignment of the invention to: Transmeta Corporation on May 24, 1999.
XXX A Declaration and Power of Attorney.
XXX A Verified Statement to establish Small Entity Status under 37 CFR 1.9 and 37 CFR 1.27.
XXX Return addressed stamped postcard.

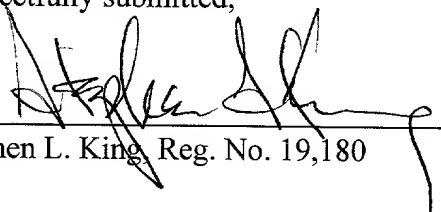
The Filing Fee has been calculated as shown below:

For:	(Col. 1)	(Col.2)	SMALL ENTITY		OTHER THAN A SMALL ENTITY	
	No. Filed	No. Extra	RATE	FEE	RATE	FEE
Basic Fee:	-	-	-	\$380.00	-	\$760.00
Total Claims:	-21	-1-	x \$9.00		x \$18.00	-0-
Indep. Claims:	-2	-0-	x \$39.00		x \$78.00	-0-
<input type="checkbox"/> Multiple Dep. Claim(s) Presented			+ \$130.00		+ \$260.00	-0-
* If the difference in (Col. 1) is less than zero, enter "0" in (Col. 2)			Total:	\$389 .00	Total:	\$.00

XXX A check in the amount of \$429.00 is enclosed to cover the filing fee and recording the assignment..

Respectfully submitted,

Date: June 14, 1999


Stephen L. King, Reg. No. 19,180

30 Sweetbay Road
Rancho Palos Verdes, California 90275
(310) 377-5073

Express Mail No.EJ738560499US

Applicant or Patentee: TRANSMETA CORPORATION Attorney's
Serial or Patent No.: _____ Docket No. TRANS11
Filed or Issued: _____

For: METHOD AND APPARATUS FOR ENHANCING SCHEDULING IN AN ADVANCED MICROPROCESSOR

VERIFIED STATEMENT (DECLARATION) CLAIMING SMALL ENTITY STATUS
37 CFR 1.9 (f) and 1.27(c) - - SMALL BUSINESS CONCERN

I hereby declare that I am an official of the small business concern empowered to act on behalf of the concern identified below:

NAME OF CONCERN: TRANSMETA CORPORATION
ADDRESS OF CONCERN: 3940 FREEDOM CIRCLE, SANTA CLARA, CALIFORNIA 95054

I hereby declare that the above identified small business concern qualifies as a small business concern as defined in 13 CFR 121.3-18, and reproduced in 37 CFR 1.9(d), for purposes of paying reduced fees under Section 41(a) and (b) of Title 35, United States Code, in that the number of employees of the concern, including those of its affiliates, does not exceed 500 persons. For purposes of this statement, (1) the number of employees of the business concern is the average over the previous fiscal year of the concern of the persons employed on a full-time, part-time or temporary basis during each of the pay periods of the fiscal year, and (2) concerns are affiliates of each other when either, directly or indirectly, one concern controls or has the power to control the other, or a third party or parties controls or has the power to control both.

I hereby certify that to the best of my knowledge and belief rights under contract or law have been conveyed to and remain with the small business concern identified above with regard to the invention entitled METHOD AND APPARATUS FOR ENHANCING SCHEDULING IN AN ADVANCED MICROPROCESSOR.

by inventor(s) Guillermo J. Rozas, Godfrey P. D'Souza, and Charles R. Price, and Paul S. Serris,
described in

[X] the specification being filed herewith,

and I have reviewed the document that evidences the conveyance of those rights. That document

[X] is being filed herewith.

If the rights held by the above-identified small business concern are not exclusive, each individual, concern or organization having rights to the invention is listed below and **no rights to the invention are held by any person, other than the inventor, who could not qualify as a small business concern under 37 CFR 1.9(d) or by any concern which would not qualify as a small business concern under 347 CFR 1.9(d) or a non-profit organization under 37 CFR 1.9(e).** NOTE: Separate verified statements are required from each named person, concern or organization having rights to the invention averring to their status as small entities. (37 CFR 1.27)

NAME: _____
ADDRESS: _____

[] Individual [] Small Business Concern [] Non-Profit Organization

NAME: _____
ADDRESS: _____

[] Individual [] Small Business Concern [] Non-Profit Organization

I acknowledge the duty to file, in this application or patent, notification of any change in status resulting in loss of entitlement to small entity status prior to paying, or at the time of paying, the earliest of the issue fee or any maintenance fee due after the date on which status as a small entity is no longer appropriate. (37 CFR 1.28(b)).

I hereby declare that all statements made herein of my own knowledge are true and that all statements made on information and belief are believed to be true; and further that these statements were made with the knowledge that willful false statements and the like so made are punishable by fine or imprisonment, or both, under Section 1001 of Title 18 of the United States Code, and that such willful false statements may jeopardize the validity of the application, any patent issuing thereon, or any patent to which this verified statement is directed.

NAME OF PERSON SIGNING: DAVID DITZEL

TITLE OF PERSON OTHER THAN OWNER: CHIEF EXECUTIVE OFFICER

ADDRESS OF PERSON SIGNING: 3940 FREEDOM WAY, SANTA CLARA, CALIFORNIA 95054

SIGNATURE:  DATE: June 7, 1999

METHOD AND APPARATUS FOR ENHANCING SCHEDULING IN AN ADVANCED MICROPROCESSOR

BACKGROUND OF THE INVENTION

Field Of The Invention

5 This invention relates to computer systems and, more particularly, to methods and apparatus for accelerating the reordering of instructions in an improved microprocessor.

History Of The Prior Art

10 Recently, a new microprocessor was developed which combines a simple but very fast host processor (called "morph host") and software (called "code morphing software") to execute application programs designed for a processor different than the morph host processor at a rate which cannot be attained by the processor for which the programs were designed (the target processor). The morph host processor executes the code morphing software to translate the application programs into morph host processor instructions which accomplish the purpose of the original target software. As the target instructions are translated, they are both executed and stored in a translation buffer where they may be accessed without further translation. Although the initial translation and
15 execution of a program is slow, once translated, many of the steps normally required to execute a program in hardware are eliminated.
20

In order to be able to execute programs designed for other processors at a rapid rate, the morph host processor includes a number of hardware enhancements. One of these enhancements is a gated store buffer which

resides between the host processor and the translation buffer. A second enhancement is a set of host registers which store state of the target machine at the beginning of any sequence of target instructions being translated. Sequences of target instructions spanning known states of the target processor are translated into morph host instructions and placed in the translation buffer awaiting execution. If the translated instructions execute without raising an exception, the target state at the beginning of the sequence of instructions is updated to the target state at the point at which the sequence completed.

If an exception occurs during the execution of the sequence of host instructions which have been translated, the processing stops; and the entire operation may be returned or rolled back to the beginning of the sequence of target instructions at which known state of the target machine exists. This allows very rapid and accurate handling of exceptions while dynamically translating and executing instructions, a result which had never been accomplished by the prior art.

Additional speed is attained in running the new microprocessor by a scheduler which is part of the code morphing software. The scheduler reorders and reschedules the instructions as they are being translated from a naive order produced by raw translation into an order which produces the same result but allows faster execution. A scheduler attempts to place certain instructions ahead of other instructions or to run instructions together so that the execution of the rescheduled software takes less time. Schedulers function with a number of constraints the most basic of which is that the rescheduled program must still produce the same ultimate results as the original program.

As an example, there are sequences of instructions in programs which must be carried out without interruption in order for the sequences to produce the correct results. A scheduler cannot interfere with such sequences without interfering with the results produced. Many
5 processors provide hardware interlocks to assure that such sequences are, in fact, run without interruption. The need to protect such sequences of instructions poses special constraints for processors without hardware interlocks such as the advanced morph host processor being discussed. Software must somehow be aware of such sequences
10 and assure that they are run without interruption.

Control dependencies are another traditional constraint on reordering which a scheduler faces. Control dependencies relate to branch instructions; a scheduler must assure that reordering of instructions which occur before and after a branch do not cause the program to run
15 incorrectly.

Other dependencies affect the reordering of loads with respect to stores. For example, if updated data is to be stored to a memory address and then manipulated in a register operation, the data at the address should not be kept in a register at the time the store occurs or the data in the
20 register may be stale.

All of these constraints cause a typical scheduler to function very conservatively and, consequently, to produce slower code.

A traditional scheduler does its best to determine those instructions which depend on one another in order to accomplish reordering. The
25 usual scheduler can determine that some operations depend on other

operations in some way and that some operations do not depend on other operations in any way, but it cannot determine anything with regard to other operations. Such a scheduler treats those operations which depend on other operations conservatively by ordering them in the normal naive order in which they originated. Such a scheduler reorders operations which do not depend on other operations at all in the manner it desires. Finally, it treats all operations about which it cannot make a determination regarding dependencies as though they depended on one another and handles them conservatively and slowly.

It is desirable to provide circuitry and software for enabling a scheduler of an advanced processor to generate code which executes at an accelerated speed.

Summary Of The Invention

The present invention is realized by apparatus and a method for causing scheduler software to produce code which executes more rapidly by ignoring some of the normal constraints placed on scheduling operations and simply scheduling certain instructions to run as fast as possible, raising an exception if the scheduling violates a scheduling constraint, and determining steps to be taken for each set of instructions about which an exception is raised.

These and other objects and features of the invention will be better understood by reference to the detailed description which follows taken together with the drawings in which like elements are referred to by like designations throughout the several views.

Brief Description Of The Drawings

Figure 1 is a diagram illustrating a new microprocessor which may utilize the present invention.

Figure 2 is a block diagram of hardware for implementing the new microprocessor of Figure 1.

Figure 3 is a flow chart illustrating a main processing loop of the new processor of Figure 1.

Figure 4 is a block diagram illustrating a portion of the new processor.

Figure 5 is a block diagram illustrating another portion of the new processor.

Figure 6 is a flow chart illustrating the operation of scheduler software designed in accordance with the invention.

Figure 7 is a block diagram illustrating one embodiment of circuitry for practicing the present invention.

Detailed Description

Figure 1 illustrates a new microprocessor 10 which combines an enhanced hardware processing portion (referred to as a “morph host”) which is much simpler than state of the art microprocessors and an emulating software portion (referred to as “code morphing software”).

The two portions function together to carry out the operations normally accomplished by hardware alone in an advanced microprocessor. The new microprocessor 10 is faster than microprocessors of the prior art, is

capable of running all of the software for all of the operating systems which may be run by a large number of families of prior art microprocessors, yet is less expensive than prior art microprocessors.

5 The microprocessor 10 includes a morph host processor 11 designed to execute code morphing software 12 for executing application programs designed for a different target processor. The morph host 11 includes hardware enhancements especially adapted to allow the acceleration techniques provided by the code morphing software 12 to be utilized efficiently. The morph host processor includes hardware enhancements to assist in accelerating operations and in providing state of a target computer immediately when an exception or error occurs. The code morphing software includes software which, among other things, translates the instructions of a target program to morph host instructions, schedules and optimizes host instructions, and responds to exceptions and errors when necessary by rolling back execution to the last point at which execution is known to have been correct and replacing working state with correct target state at that point so that correct retranslations of target code may occur. Code morphing software also includes various processes for enhancing the speed of processing.

10

15

20 The block diagram of Figure 2 illustrates in detail exemplary hardware of a morph host 11 which implements the features discussed herein.

As is illustrated in the diagram of Figure 3 (which describes the operation of the main loop of the code morphing software 12), the code morphing software combined with the enhanced morph host translates target instructions into instructions for the morph host on the fly and caches those host instructions in a memory data structure (referred to as a

25

“translation buffer”). Once a target instruction has been translated, it may be recalled from the translation buffer and executed without the need for any of the myriad of steps required by prior art hardware microprocessors such as: determining which primitive instructions are required to implement each target instruction, addressing each primitive instruction, fetching each primitive instruction, optimizing the sequence of primitive instructions, allocating assets to each primitive instruction, reordering the primitive instructions, and executing each step of each sequence of primitive instructions involved each time each target instruction is executed.

A primary problem of prior art emulation techniques has been the inability to handle exceptions generated during the execution of a target program with good performance. Some exceptions generated in running the target application are directed to the target operating system, and the correct target state must be available at the time of any such exception for proper execution of the exception and the instructions which follow. Other exceptions can be generated by the emulator to detect particular target operations which have been replaced by some particular host function. The host processor executing the host instructions derived from the target instructions can also generate exceptions. All of these exceptions can occur either during the attempt to change target instructions into host instructions by the emulator, or when the emulating host instructions are executed by the host processor. Exceptions directed to the target operating system are especially difficult because they require a knowledge of the state of the target processor at all times.

In order to recover from these exceptions efficiently, the enhanced morph host includes a number of hardware improvements. These improvements include a gated store buffer (see Figure 5). The gated store buffer stores working memory state changes on an “uncommitted” side of a hardware “gate” and official memory state changes on a “committed” side of the hardware gate where these committed stores “drain” to main memory. A “commit” operation transfers memory stores from the uncommitted side of the gate to the committed side of the gate. If an exception occurs, a “rollback” operation discards uncommitted stores in the gated store buffer.

The hardware improvements also include a large plurality of additional processor registers (see Figure 4). In addition to allowing register renaming to lessen the problem of instructions trying to utilize the same hardware resources, the additional registers allow the maintenance of a set of host or working registers for processing the host instructions and a set of target registers to hold the official state of the target processor for which the target application was originally created. The target registers are connected to their working register equivalents through a dedicated interface that allows a commit operation to quickly transfer the content of all working registers to official target registers and allows an operation called “rollback” to quickly transfer the content of all official target registers back to their working register equivalents.

The additional official registers and the gated store buffer allow the state of memory and the state of the target registers to be updated together once one or a group of target instructions have been translated and run without error. Updates are chosen by the code morphing software to

occur on integral target instruction boundaries. If the primitive host instructions generated by a translation of a series of target instructions are run by the host processor without generating an exception, then the working memory stores and working register state generated by those instructions are transferred to official memory and to the official target registers.

On the other hand, if an exception occurs when processing the host instructions at a point which is not on a target instruction boundary, the original state in the target registers at the last update (or commit) may be recalled to the working registers and uncommitted memory stores in the gated store buffer may be discarded. Then, if the exception generated is a target exception, the target instructions causing the target exception may be retranslated one at a time and executed in serial sequence as they would be executed by a target microprocessor. As each target instruction is correctly executed without error, the state of the target registers may be updated; and the data in the store buffer gated to memory. Then, when the exception occurs again in running the host instructions, the correct state of the target processor is held by the target registers of the morph host and memory; and the operation may be correctly handled without delay. Each new translation generated by this corrective translating may be cached for future use as it is translated or alternatively discarded if caused by a one time or rare occurrence such as a page fault. These features combine to assist the microprocessor created by the combination of the code morphing software and the morph host to execute instructions more rapidly than processors for which the software was originally written.

In addition to simply translating the instructions, caching the translated instructions, and executing each translation whenever that set of instructions needs to be executed, the code morphing software also reorders, optimizes, and reschedules the different translations. One
5 optimizing process links the various sequences of translated host instructions to one another as the probable branches to be taken become apparent during execution. Eventually, the main loop references in the branch instructions of host instructions are almost completely eliminated. When this condition is reached, the time required to fetch
10 target instructions, decode target instructions, fetch the primitive instructions which make up the target instructions, optimize those primitive operations, reorder the primitive operations, and reschedule those primitive operations before running any host instruction is eliminated. Thus, the work required to run any set of target instructions
15 using the improved microprocessor is drastically reduced.

As pointed out above, the operation of reordering utilizes a scheduler which attempts to choose a better order for instructions to execute when presented correctly but naively ordered instructions. A problem with schedulers is that they function with a number of constraints. The most
20 basic constraint is that the program must still produce the same ultimate result when executing as did the original sequence of instructions. All of these constraints cause a typical scheduler to function very conservatively and, consequently, produce code which executes slowly.

25 For example, to ensure that a correct result is produced, the typical scheduler operates upon a deterministic basis to select those

instructions which have no dependencies, those instructions which have dependencies, and those instructions about which the presence of dependencies is unknown. Those instructions which have dependencies and those about which the presence of dependencies is unknown are all
5 treated as though dependencies exist and are not reordered. Only those instructions known to be without dependencies are reordered. Following these guidelines, schedulers produce code which executes slowly.

Another constraint relates to a particular embodiment of the morph host processor. One embodiment of the morph host processor is a processor
10 which is designed to function rapidly by eliminating specialized circuitry which slows operations. This embodiment of the morph host processor is designed without any hardware locking mechanism. A hardware locking mechanism is a circuit intended to assure that all steps in a particular sequence of instructions are executed without being
15 interrupted. The lack of a locking mechanism requires a scheduler to function very strictly to assure that all steps in such sequences are handled in the originally translated order without any reordering so that the processor will produce the correct result from the sequence.

The scheduler of the present invention is a software portion of the code
20 morphing software. Unlike hardware schedulers of the prior art, the software scheduler uses speculative techniques in reordering instructions. The scheduler speculates that for certain operations the fastest possible operation is desired and reorders instructions to accomplish this result. Hardware is provided in the morph host to raise
25 an exception if the speculation chosen is incorrect. In most cases, the speculation is correct, so the overall result is much faster operation.

However, if the speculation is incorrect, then the exception typically causes the software to utilize the gated store buffer and the target registers to roll back operations to the beginning of the speculative sequence at which correct state is known.

5 In contrast to the deterministic strategy used by prior art schedulers, the scheduler of the present invention utilizes probabilistic guidelines in selecting categories of instructions for reordering. The improved scheduler selects four categories of sequences of instructions (see Figure 6) from the sequences of instructions produced by translation from a set of target instructions. These categories include sequences of instructions with no dependencies, sequences of instructions with known dependencies, sequences of instructions which probably have no dependencies, and sequences of instructions which probably have dependencies. As with the prior art, those sequences of instructions known to have no dependencies may be reordered at will by the scheduler; and those sequences of instructions with known dependencies are handled in the sequential order provided by the translator.

However, the instructions which probably have no dependencies are treated as though they in fact have no dependencies and are reordered in a manner to provide the fastest possible execution. Hardware means are provided in the morph host to detect an incorrect reordering and raise an exception if dependencies in fact exist. The scheduler cooperates with the hardware means to ensure that a check is enabled for each reordered instruction which may execute incorrectly in order to raise an exception when the sequence of operations does execute incorrectly. Such an exception allows the scheduler to ignore its previous reordering which

raised the exception and treat the sequence conservatively or in some other more appropriate manner.

Those instructions which probably have dependencies, on the other hand, may be handled either aggressively or conservatively. If handled aggressively, they are treated as are those instructions which probably have no dependencies. They are reordered in a manner to provide the fastest possible execution and use the hardware means provided in the morph host to detect and raise an exception if incorrect reordering has taken place. If handled conservatively, they are handled in the sequential order provided by the translator. Normally, the conservative treatment will provide faster processing because the raising of a large number of exceptions noticeably slows execution speed.

In one embodiment of the present invention, circuitry such as that illustrated in Figure 7 is added to the host processor. This circuitry is utilized to store a memory address accessed by an instruction which has been reordered by the scheduler using a special "load and protect" or "store and protect" operation. Such a "load and protect" or "store and protect" operation is used whenever an instruction has been reordered and has the effect of placing the memory address accessed by the reordered instruction in one of a plurality of registers 71 of the morph host designated for use as protection registers. In one embodiment, eight protection registers 71 are provided. The "load and protect" or "store and protect" instruction indicates the particular protection register which is to be used for the operation.

Although the term “memory address” is used throughout this specification in explaining the load and protect and store and protect instructions, the term is used as a reference to a number of possible arrangements for determining a memory region which is protected. The
5 term memory address is used to mean a descriptor of a memory region which is being protected. For example, in a system in which memory is byte addressable, one embodiment of the invention uses a starting memory address and a number of bits equal to the number of bytes within the address region to indicate the protected condition of each of
10 those bytes. Another embodiment with similar addressing utilizes a starting memory address and a length, while a third embodiment utilizes individual byte addresses and individual comparators for each byte address.

In an exemplary operation a sequence of instructions includes in order a
15 first store instruction STORE1, a second store instruction STORE2, and a load instruction LOAD1. The scheduler decides to reorder the instructions to place the load instruction first, the second store instruction second, and the first store instruction third in the reordered sequence on the assumption that the reordering probably will not cause
20 an incorrect operation. To do this, the scheduler uses the “load and protect” operation to place the load data in one of the general registers 72 and the address of the memory position from which the load data was derived in the protection register 71 designated by the instruction. Because the software scheduler understands which instructions need be
25 checked to determine whether the reordering has caused an error, the scheduler places an indication (e.g., a bit in a bitmask) in those next

instructions which may be affected by the reordering (in this case, the STORE1 and STORE2 instructions ahead of which the load was placed) to indicate the particular protection register holding the protected memory address. The presence of this indication in a particular position
5 (one of eight bits if eight protection registers are used for the trapping function) indicates that the execution of the instruction depends on whether the address in which each store is to be placed by the store instruction overlaps the memory address held in the indicated protection register 71.

Similarly, the scheduler uses a “store and protect” operation to store the data of the STORE2 instruction in memory and to place the address of the memory position to which the data was stored in the protection register 71 designated by the store and protect instruction. The scheduler also places an indication in the bitmasks of each of the
10 instructions which may be affected by the reordering (in this case, the STORE1 instruction only) to indicate the particular protection register holding this protected memory address. Finally, the scheduler uses a normal store instruction for the last STORE1 instruction.
15

As the sequence of instructions is executed, the host hardware uses
20 comparator circuitry 73 to determine for each of these three instructions if the memory address of the instruction overlaps any portion of the data at a memory address stored in one of the protection registers 71 and, if so, raises an exception. Thus, the LOAD1 operation (which has become a load and protect) writes its memory to a protection register 71 but
25 does not check any protection registers since none have been designated by indicators being set. The STORE2 operation (which has become a

store and protect) writes its memory position to a different protection register 71 and checks the protection register 71 used for the LOAD1 instruction to determine overlap between their memory positions. Finally, the STORE1 operation (which remains a simple store although augmented by the protection register indicators) checks the protection registers for each of the LOAD1 and STORE2 instructions for overlap between its memory address and the memory addresses of the LOAD1 and STORE2 instructions. In the case of the first and third embodiments described above, the comparison allows protection to be precisely applied to the byte level.

Any exception causes the code morphing software to determine the steps to be taken in response to the exception. Typically, the code morphing software causes the execution of the reordered sequence of instructions to be aborted and the host to revert to the state of the target processor at the beginning of the sequence of instructions so that the sequence of instructions may be reprocessed conservatively. If the addresses are not the same (indicating in the example that the store instruction does not access the protected memory address), the execution of the reordered sequence of instructions proceeds at the accelerated pace provided by the reordering.

In order to implement communications between the host processor and the scheduler, the load and store instructions utilized by the morph host have been modified. In one embodiment, these instructions are completely replaced by the “load and protect” and “store and protect” instructions. Each “load and protect” and each “store and protect” instruction includes a bitmask (e.g., eight bits corresponding to eight

protection registers) the bits of which are used as flags to indicate a particular protection register in which to look for the memory address of the reordered or aliased instruction. Each of these bits designates one of the available protection registers in which the memory addresses are to be stored for the hardware to check. With this bitmask, the specific protection register designated to store the memory address when the instruction is reordered may be checked before the subsequent instruction which may be affected by the reordering is executed. The “load and protect” and “store and protect” instructions may be used in place of the normal load and store instructions, respectively, because when no bits of the bitmask are set, no checking will occur. In such a case, “load and protect” and a “store and protect” operations are identical to load and store operations. It should also be noted that the ability to associate protection registers with particular general registers holding memory data allows efficient use of a small number of protection registers.

The host processor of the present invention also includes an additional register called an enable protection register 74 to store the positions of the protection registers which contain valid memory addresses related to reordered instructions. The bits indicating particular protection registers are set to indicate the protection registers using the indication provided by the “load and protect” or “store and protect” instructions. In one embodiment, the bits of the enable protection register are cleared whenever a commit operation occurs indicating that a sequence of translated and reordered instructions has executed without raising a reordering exception. Since reordering only takes place in sequences of

instructions all of which occur between two commit points, this allows the reordering operation to utilize all of the protection registers allotted to reordering for each newly translated sequence of instructions.

An additional advantage of the new invention is that the “store and protect” operation allows reordering of stores with respect to one another. With the present invention, this may be accomplished by storing data to a memory position and protecting the address of the memory position in a protection register. When the subsequent store which may be affected by the reordering occurs, its bitmask indicates the protection register which the hardware should check for a memory address to determine if an exception should be raised or if the reordering of stores has been accomplished correctly.

In one embodiment of the new microprocessor, a circuit arrangement has been provided that allows memory data used quite often in the execution of an operation to be replicated (or “aliased”) in an execution unit register in order to eliminate the time required to fetch the data from or store the data to memory. For example, if data in memory is reused frequently during the execution of one or more code sequences, the data must typically be retrieved from memory and loaded to a register in an execution unit each time the data is used. To reduce the time required by such frequent memory accesses, the data may instead be loaded once from memory to an execution unit register at the beginning of the code sequence and the register designated to function in place of the memory space during the period in which the code sequence continues. Once this has been accomplished, each of the load operations which would normally involve loading data to a register from the designated memory

address becomes instead a simple register-to-register copy operation which proceeds at a much faster pace; and even those copy operations may frequently be eliminated by further optimization.

Similarly, execution of a code sequence often requires that data be written to a memory address frequently during the execution of a code sequence. To reduce the time required by such frequent memory stores to the same address, each time the data is to be written to the memory address, it may be transferred to an execution unit register which is designated to function in place of the memory space during the period in which the code sequence is continuing. Once an execution unit register has been designated, each change to the data requires only a simple register-to-register transfer operation which proceeds much faster than storing to a memory address.

The operation of the aliasing circuitry is described in U. S. patent application serial no. 08/721,698, filed September 26, 1996, entitled Method and Apparatus for Aliasing Memory Data in an Advanced Microprocessor, M. Wing et al, and assigned to the assignee of the present invention.

A second embodiment of the present invention for accelerating reordered operations utilizes certain additional hardware so that the same hardware may be used both for reordering and for aliasing of memory addresses in the manner described in the above-mentioned patent application. It should be noted that reordering instructions typically occurs in intervals between adjacent commit operations while aliasing memory data in an execution unit register typically remains in effect for

much longer periods. In this second embodiment, a second “persistent” register 76 is added to allow long term or persistent protection to be utilized along with the short term protection provided for reordering by the enable protection register 74. The second persistent register 76 is used in the same manner as the register 74 but to record only those protection registers in which memory addresses should be maintained for periods longer than between adjacent commit operations.

For example, if it is desired to alias a memory address and store the data in a host register to be used for some long period (e.g., during a loop), the indication of which protection register is holding the address of the long term aliasing operation is copied from the instruction and is placed in both the enable protection register 74 and the second persistent register 76. Presuming that the sequence of reordered instructions executes without raising an exception thereby allowing a first commit operation to occur, then the enable protection register is cleared. In this manner, the short term flags indicating the protection registers holding the addresses of reordered instructions to check are eliminated at each commit. After the enable protection register is cleared at the commit, the contents of the second persistent register are written into the enable protection register. Since the data in the persistent register indicating which protection registers are being used for long term aliasing is written to the enable protection register, the indications of the protection registers used for long term aliasing are not affected by the commit operation. By writing the contents of the persistent register to the enable protection register at each commit, protection is effectively continued for the next sequence of instructions and, ultimately, until the second register is

finally cleared when the data is no longer needed for the aliasing operation.

In addition to the second persistent register 76, a shadow register 78 is kept which stores the information also kept in the persistent register.

5 The shadow register is used during commit and rollback operations.

When a commit occurs, the data in the persistent register 76 is copied in the manner discussed above to the enable protection register 74. The same data is also copied on a commit to the register 78 shadowing the persistent register so that the shadow register contains the settings for the persistent register at the start of the next sequence of instructions which may be reordered. If an exception occurs during the execution of the next sequence of instructions and a rollback operation is necessary, then the contents of the shadow register are copied to both the enable protection register and the persistent register. This places the same indications in the enable protection and persistent registers as was in those registers before the execution of the sequence of instructions began thereby assuring correct state for the more conservative execution which follows.

Additional advantages are provided by the arrangement of the present invention. While the addition of the persistent register 76 allows the use of the same hardware for both enhancing the ability to reorder during the short term (between commits) and to maintain aliased memory data for long periods in execution unit registers thereby eliminating memory accessing redundancies, it may also be used to eliminate other types of redundancies which occur between commit operations. For example, it is possible that two loads from the same memory address may occur

during a sequence of instructions. If this happens and there are no intervening stores to the memory address, then the second load may be simply ignored; and the data placed in a register by the first memory access used without change in place of the second load operation.

5 However, if a store intervenes between the loads, it is necessary to determine whether the store occurred to the memory address from which the second access is to occur. Thus, prior art optimization techniques would not be able to eliminate the second load if a store intervened between the loads.

10 In order to shorten the operation, the present invention may be used to advantage. If the first load is changed to a "load and protect" operation with the memory address stored in a protection register and the store instruction receives a flag to indicate the particular protection register to check, then the second load may be eliminated and the data stored by
15 the "load and protect" operation used for the second load. If the store instruction attempts to access the protected memory address, the flag indicating the protection register to check will cause a comparison to take place before the store access occurs. This will generate an exception, and a rollback to the last commit point at which correct target
20 state exists will occur. The scheduler may then provide the appropriate instruction sequence including the second load operation and the sequence may be reexecuted

25 Similarly, if a sequence of instructions between two commit operations includes two stores to the same memory address, the first store may be eliminated if no load from the memory address has occurred between the stores. However, if data from the memory address has been used for a

load in the interim, then the first store may not be eliminated. Using the present invention, the first store to the memory address may be eliminated if the load instruction is made a “load and protect.” Then the second store receives the protection register indication from the “load and protect” to check the memory address of the access. If the load is from a different address, then the second store may proceed correctly. If the load is from the same address, then the attempt to access memory for the second store will generate an exception which will roll back the operation to the last commit point. From that point the scheduler may reschedule the instructions to include both store operations and reexecute the sequence.

Although the present invention has been described in terms of a preferred embodiment, it will be appreciated that various modifications and alterations might be made by those skilled in the art without departing from the spirit and scope of the invention. For example, although the invention has been the embodiment described has been designed to function with a particular family of processors, it should be understood that the invention applies just as well to programs designed for other processor architectures, and programs. The invention should therefore be measured in terms of the claims which follow.

What Is Claimed Is:

1 Claim 1. A method for causing scheduler software to produce code
2 which executes rapidly including the steps of:
3 reordering a sequence of instructions to run as fast as possible even
4 though the reordered sequence may generate an exception,
5 raising an exception if the reordered sequence of instructions violates a
6 scheduling constraint, and
7 determining steps to be taken for correctly executing each set of
8 instructions about which an exception is raised.

1 Claim 2. A method as claimed in claim 1 in which the step of raising
2 an exception if the reordered sequence of instructions violates a
3 scheduling constraint includes a step of detecting exceptions caused by
4 reordered instructions in the reordered sequence.

1 Claim 3. A method as claimed in claim 2 in which the step of
2 detecting exceptions caused by reordered instructions in the reordered
3 sequence includes;
4 remembering an instruction which has been placed out of order in the
5 sequence, and
6 checking instructions in the sequence with respect to which the
7 remembered instruction has been reordered to determine if an incorrect
8 result is produced by the sequence of instructions.

1 Claim 4. A method as claimed in claim 2 in which the step of
2 remembering an instruction which has been placed out of order in the
3 sequence includes storing a memory address accessed by the
4 instruction, and
5 the step of checking instructions in the sequence with respect to which
6 the remembered instruction has been reordered to determine if an
7 incorrect result is produced by the sequence of instructions includes
8 comparing the stored memory address with memory addresses accessed
9 by instructions against which it is checked.

1 Claim 5. A method as claimed in claim 4 in which the memory
2 address of the data accessed by the instruction is stored in a protection
3 register.

1 Claim 6. A method as claimed in claim 4 including a further step of
2 remembering where all remembered instructions are held.

1 Claim 7. A method as claimed in claim 1 in which it is probable that
2 the reordered sequence will generate an exception.

1 Claim 8. A method as claimed in claim 7 in which the step of raising
2 an exception if the reordered sequence of instructions violates a
3 scheduling constraint includes a step of detecting exceptions caused by
4 reordered instructions in the reordered sequence.

1 Claim 9. A method as claimed in claim 8 in which the step of
2 detecting exceptions caused by reordered instructions in the reordered
3 sequence includes;

4 remembering an instruction which has been placed out of order in the
5 sequence, and
6 checking instructions in the sequence with respect to which the
7 remembered instruction has been reordered to determine if an incorrect
8 result is produced by the sequence of instructions.

1 Claim 10. A method as claimed in claim 8 in which the step of
2 remembering an instruction which has been placed out of order in the
3 sequence includes storing a memory address accessed by the
4 instruction, and
5 the step of checking instructions in the sequence with respect to which
6 the remembered instruction has been reordered to determine if an
7 incorrect result is produced by the sequence of instructions includes
8 comparing the stored memory address with memory addresses accessed
9 by instructions against which it is checked.

1 Claim 11. A method as claimed in claim 10 in which the memory
2 address accessed by the instruction is stored in a protection register.

1 Claim 12. A method as claimed in claim 10 including a further step of
2 remembering where all remembered instructions are held.

1 Claim 13. A system for controlling the reordering of instructions to
2 produce code which executes rapidly comprising:
3 means for reordering a sequence of instructions to run as fast as possible
4 even though the reordered sequence may generate an exception,

5 means for raising an exception if the reordered sequence of instructions
6 violates a scheduling constraint, and
7 means determining steps to be taken for correctly executing each set of
8 instructions about which an exception is raised.

1 Claim 14. A system as claimed in claim 13 in which the means for
2 raising an exception if the reordered sequence of instructions violates a
3 scheduling constraint includes means for detecting exceptions caused by
4 reordered instructions in the reordered sequence.

1 Claim 15. A system as claimed in claim 14 in which the means for
2 detecting exceptions caused by reordered instructions in the reordered
3 sequence comprises:

4 means for identifying an instruction which has been placed out of order
5 in the sequence, and

6 means for checking instructions in the sequence with respect to which
7 the identified instruction has been reordered to determine if an incorrect
8 result is produced by the sequence of instructions.

1 Claim 16. A system as claimed in claim 14 in which the means for
2 identifying an instruction which has been placed out of order in the
3 sequence includes means for storing a memory address accessed by the
4 instruction, and

5 the means for checking instructions in the sequence with respect to
6 which the identified instruction has been reordered to determine if an

7 incorrect result is produced by the sequence of instructions includes
8 means for comparing the stored memory address with memory addresses
9 accessed by instructions against which it is checked.

1 Claim 17. A system as claimed in claim 16 in which the means for
2 storing the memory address of the data accessed by the instruction is a
3 protection register.

1 Claim 18. A system as claimed in claim 16 further comprising means
2 for storing indicators for the memory addresses accessed by identified
3 instructions.

1 Claim 19. A system as claimed in claim 16 further comprising means
2 for storing indicators for the memory addresses accessed by identified
3 instructions which exist during a sequence of instructions, and
4 means for storing indicators for the addresses accessed by identified
5 instructions which persist beyond a sequence of instructions.

1 Claim 20. A system as claimed in claim 19 in which the means for
2 storing indicators for the memory addresses accessed by identified
3 instructions which exist during a sequence of instructions is a first
4 register storing indications of valid memory addresses accessed by
5 identified instructions,
6 the means for storing indicators for the memory addresses accessed by
7 identified instructions which persist a sequence of instructions is a
8 second register storing indications of valid memory addresses accessed
9 by identified instructions; and

10 further including a third register for storing indicators for the memory
11 addresses accessed by identified instructions, and
12 means for transferring indications from the second register to the first
13 and third registers when a sequence of instructions is executed without
14 an exception, and
15 means for transferring indications from the third register to the first and
16 second registers when an exception is generated during a sequence of
17 instructions being executed.

1 Claim 21. A system as claimed in Claim 13 further comprising means
2 for replicating memory data in an execution unit register.

Abstract of the Disclosure

Apparatus and a method for causing scheduler software to produce code which executes more rapidly by ignoring some of the normal constraints placed on its scheduling operations and simply scheduling certain

- 5 instructions to run as fast as possible, raising an exception if the scheduling violates a scheduling constraint, and determining steps to be taken for correctly executing each set of instructions about which an exception is raised.

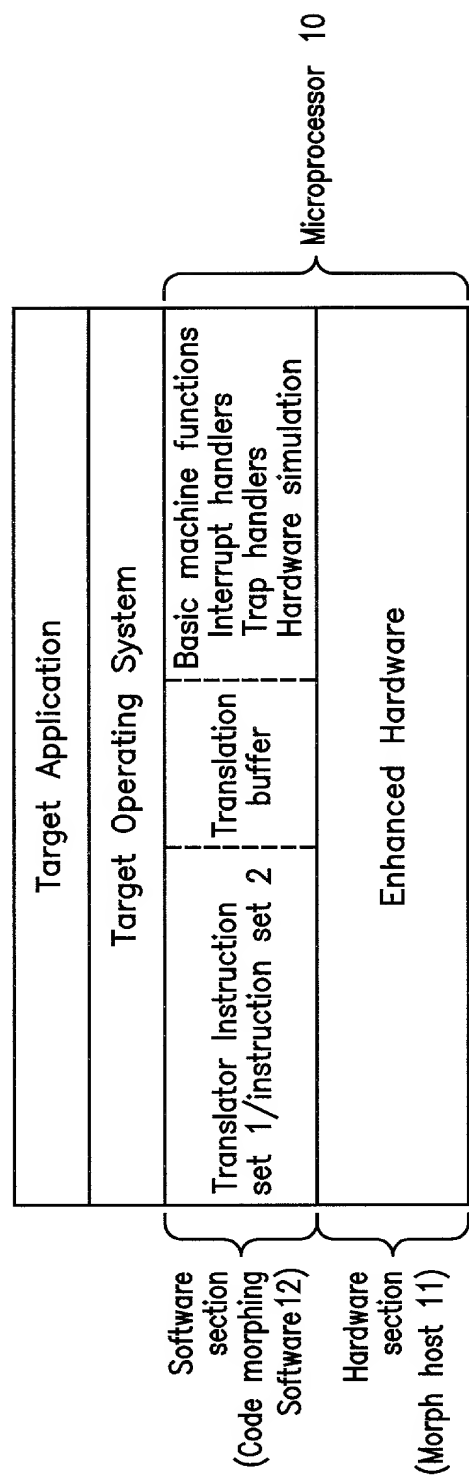


Fig. 1

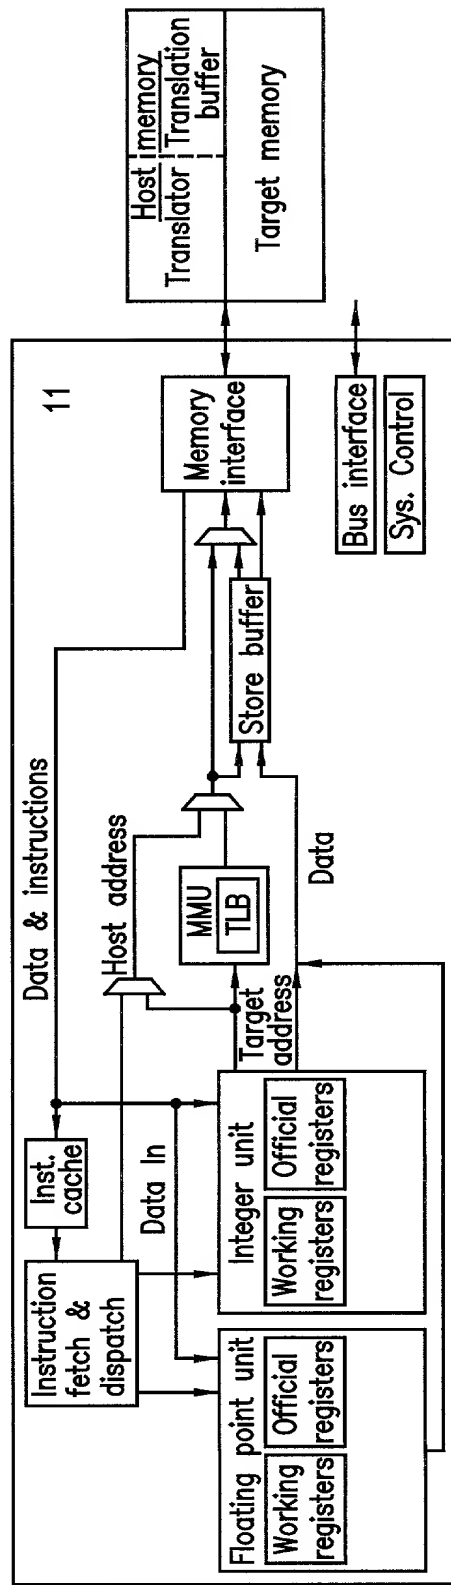


Fig. 2

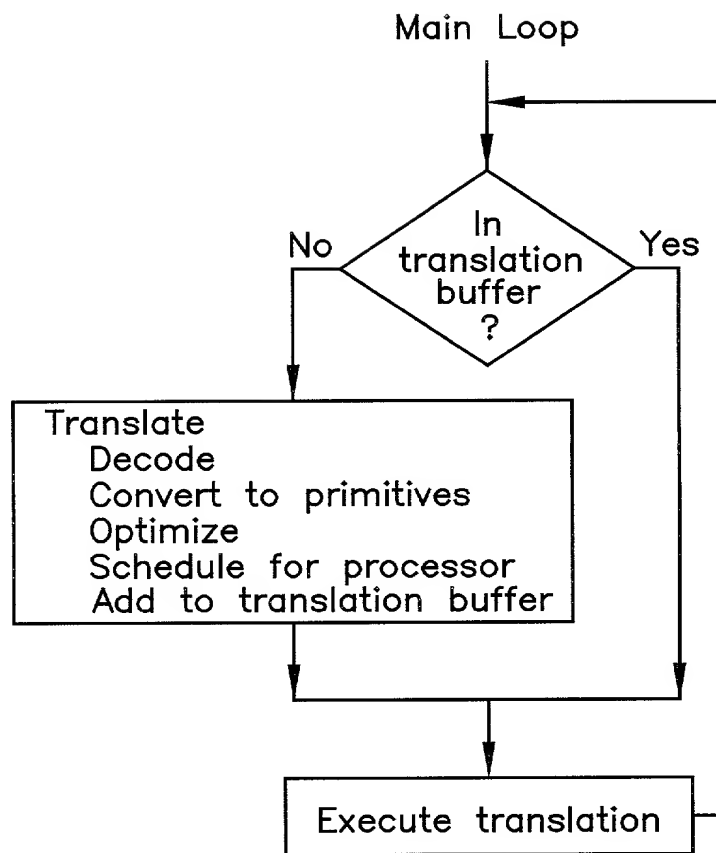


Fig. 3

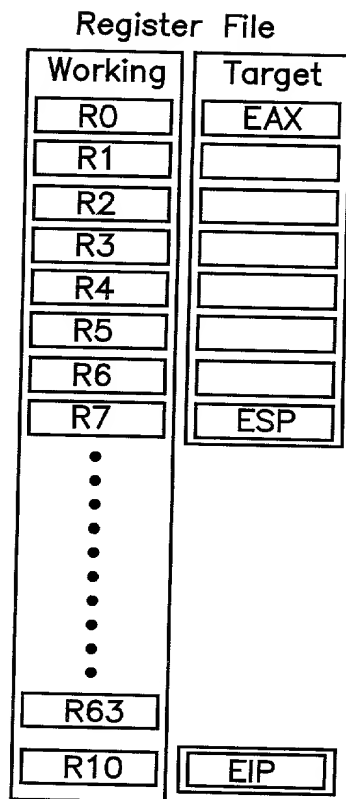


Fig. 4

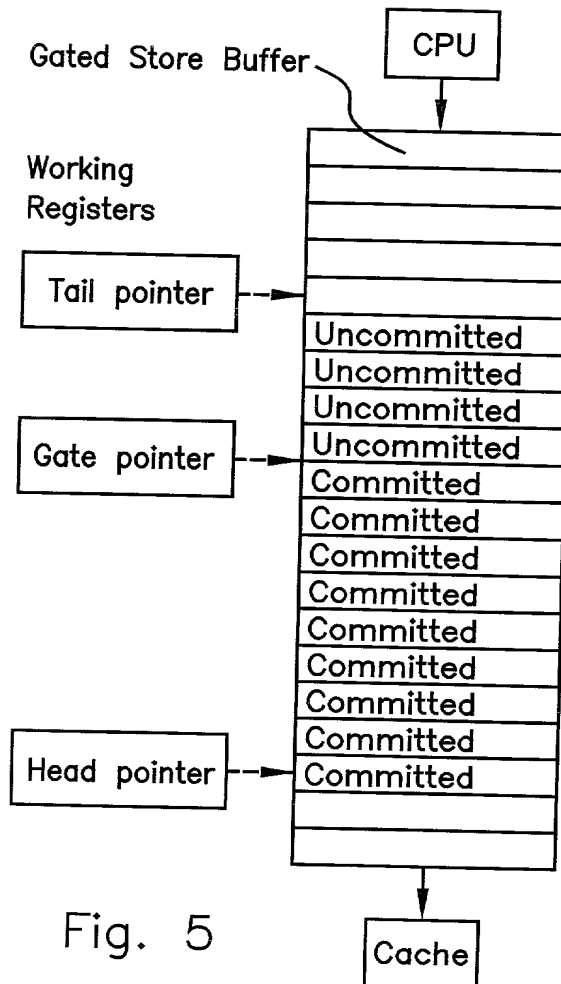
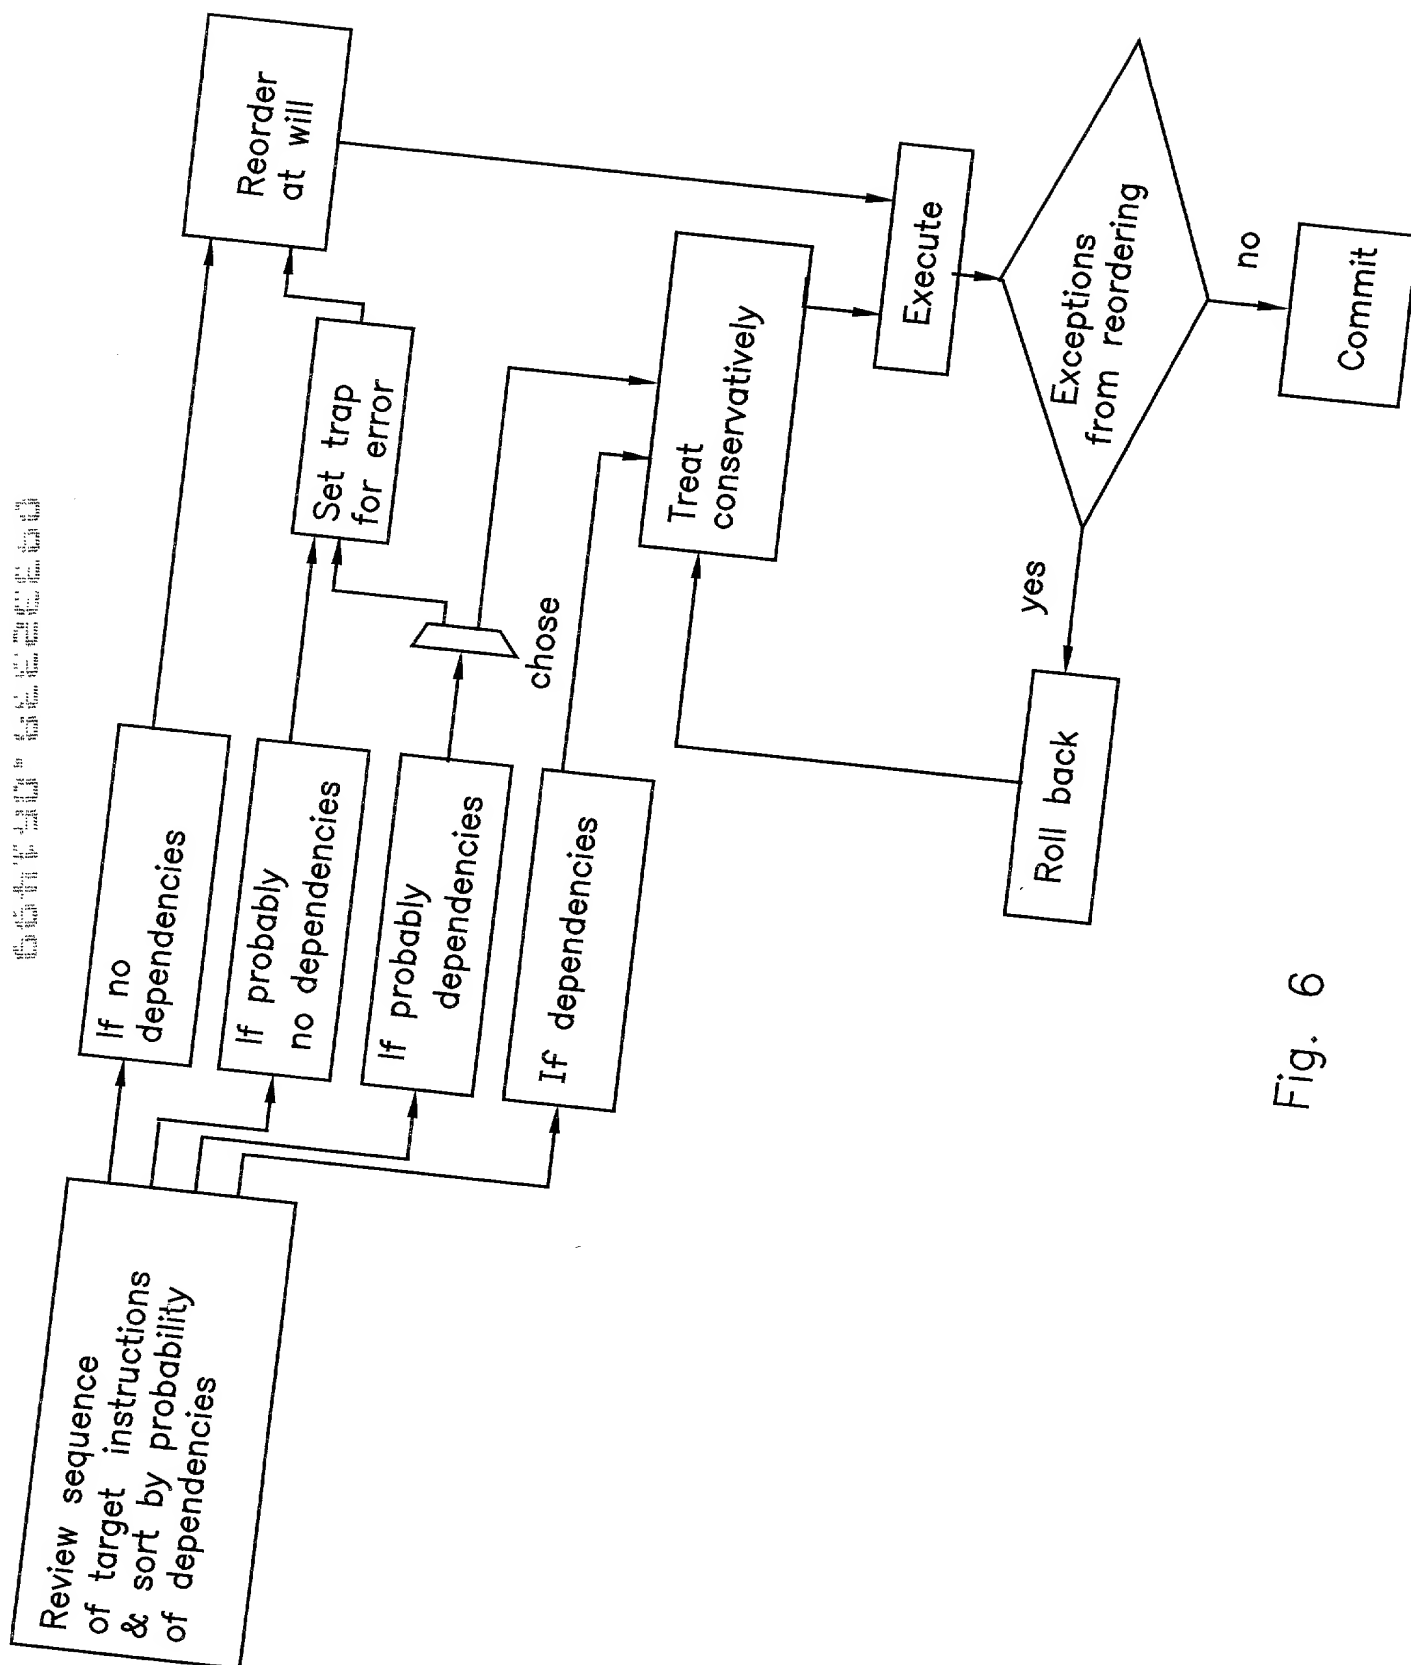


Fig. 5



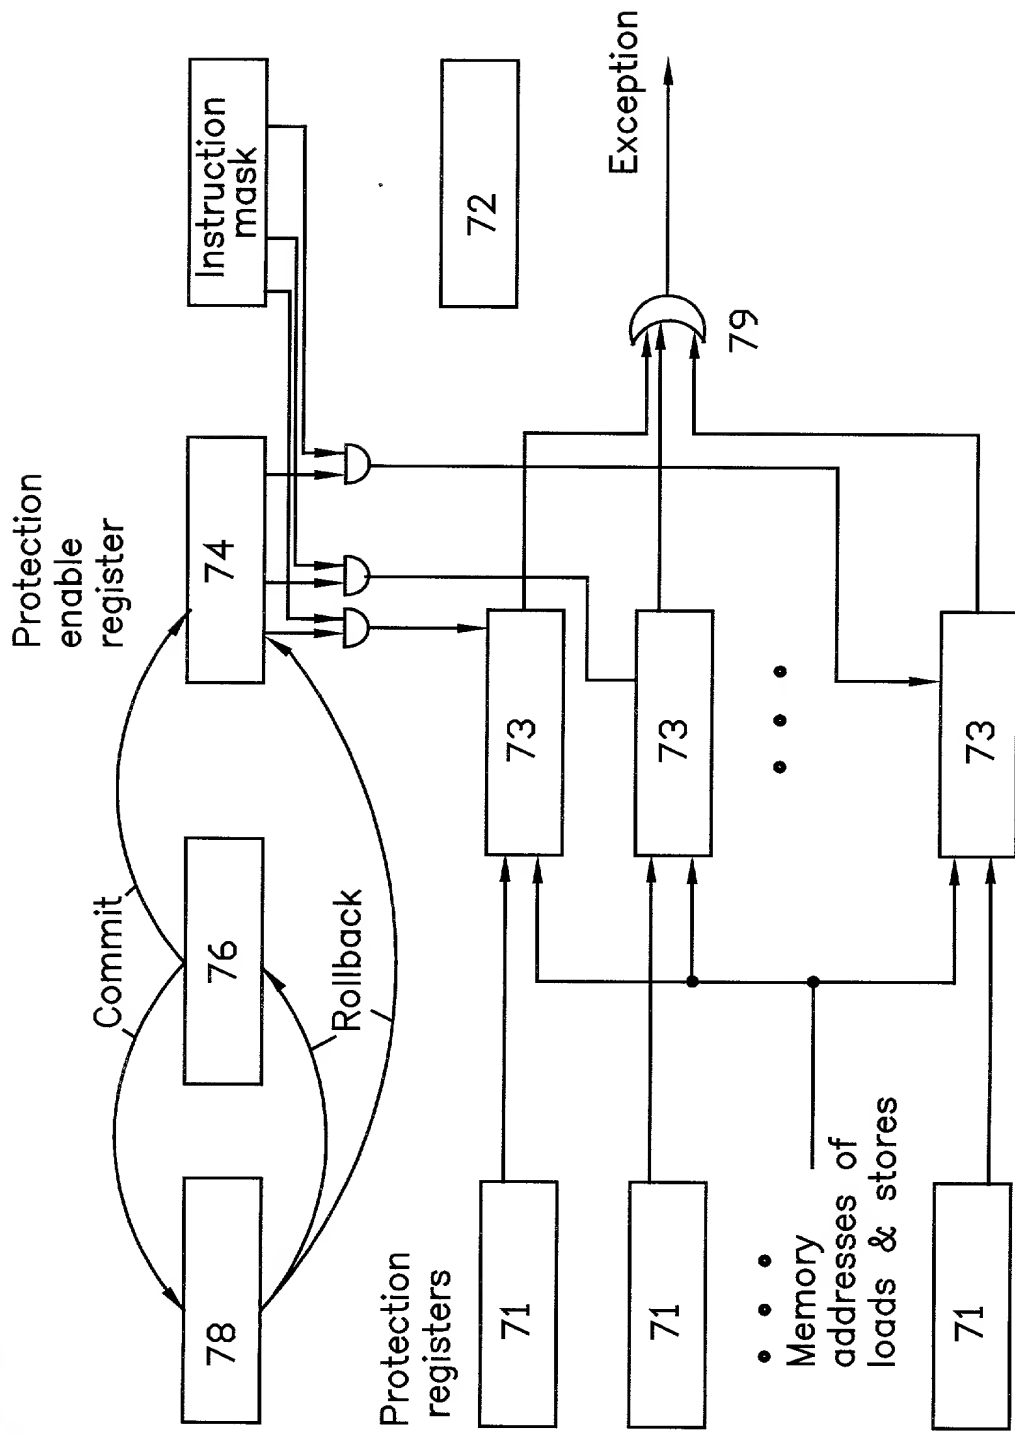


Figure 7

DECLARATION AND POWER OF ATTORNEY FOR PATENT APPLICATION

As a below named inventor, I hereby declare that:

My residence, post office address and citizenship are as stated below, next to my name.

I believe I am the original, first, and sole inventor (if only one name is listed below) or an original, first, and joint inventor (if plural names are listed below) of the subject matter which is claimed and for which a patent is sought on the invention entitled

METHOD AND APPARATUS FOR ENHANCING SCHEDULING IN AN ADVANCED MICROPROCESSOR

the specification of which is attached hereto.

I hereby state that I have reviewed and understand the contents of the above-identified specification, including the claim(s), as amended by any amendment referred to above. I do not know and do not believe that the claimed invention was ever known or used in the United States of America before my invention thereof, or patented or described in any printed publication in any country before my invention thereof or more than one year prior to this application, that the same was not in public use or on sale in the United States of America more than one year prior to this application, and that the invention has not been patented or made the subject of an inventor's certificate issued before the date of this application in any country foreign to the United States of America on an application filed by me or my legal representatives or assigns more than twelve months (for a utility patent application) or six months (for a design patent application) prior to this application.

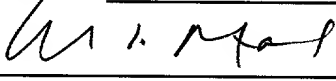
I acknowledge the duty to disclose all information known to me to be material to patentability as defined in Title 37, Code of Federal Regulations, Section 1.56.

I hereby appoint Stephen L. King, Reg. No. 19,180; with offices located at 30 Sweetbay Road, Rancho Palos Verdes, California 90275, telephone (310) 377-5073, with full power of substitution and revocation, to prosecute this application and to transact all business in the Patent and Trademark Office connected herewith.

I hereby declare that all statements made herein of my own knowledge are true and that all statements made on information and belief are believed to be true; and further that these statements were made with the knowledge that willful false statements and the like so made are punishable by fine or imprisonment, or both, under Section 1001 of

Title 18 of the United States Code and that such willful false statements may jeopardize the validity of the application or any patent issued thereon.

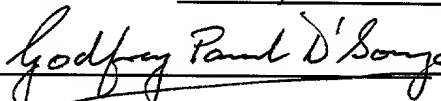
Full Name of Sole/First Inventor Guillermo J. Rozas

Inventor's Signature  Date May, 14, 1999

Residence Los Gatos, California Citizenship U.S.A.
(City, State) (Country)

Post Office Address 104 Magneson Terrace
Los Gatos, California 95032

Full Name of Second/Joint Inventor Godfrey P. D'Souza

Inventor's Signature  Date 05/14/99

Residence San Jose, California Citizenship India
(City, State) (Country)

Post Office Address 298 S. 12th Street
San Jose, California 95112

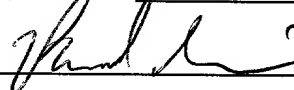
Full Name of Third/Joint Inventor Charles R. Price

Inventor's Signature  Date 19-May-1999

Residence Sunnyvale, California Citizenship U.S.A.
(City, State) (Country)

Post Office Address 455 Crescent Ave., #1
Sunnyvale, California 94087

Full Name of ^{Fourth}~~Third~~/Joint Inventor Paul S. Serris

Inventor's Signature  Date 5/24/99

Residence Sunnyvale, California Citizenship U.S.A.
(City, State) (Country)

Post Office Address 661 San Diego Ave.
Sunnyvale, California 94086